

UNIFIED CLUSTERING TREE

RELATED APPLICATIONS

[0001] This application claims priority under 35 U.S.C. § 119 based on U.S. Provisional Application No. 60/419,214 filed October 17, 2002, the disclosure of which is incorporated herein by reference.

BACKGROUND OF THE INVENTION

A. Field of the Invention

[0002] The present invention relates generally to speech recognition and, more particularly, to speech recognition models.

B. Description of Related Art

[0003] Speech has not traditionally been valued as an archival information source. As effective as the spoken word is for communicating, archiving spoken segments in a useful and easily retrievable manner has long been a difficult proposition. Although the act of recording audio is not difficult, automatically transcribing and indexing speech in an intelligent and useful manner can be difficult.

[0004] Speech recognition systems are generally based on statistical models. The models are trained on a speech signal and a corresponding transcription of the speech signal. The models “learn” how the speech signal corresponds to the transcription. Conventional models are frequently implemented based on Hidden Markov Models (HMMs).

[0005] Fig. 1 is a diagram illustrating a conventional speech recognition system. A content transcription component 102 receives an input audio stream. The content transcription component 102 converts speech in the input audio stream into text based on language and acoustic model(s) 101. Model(s) 101 are pre-trained based on a training audio stream that is expected to be similar to the run-time version of the input audio stream.

[0006] Fig. 2 is a diagram illustrating training of models 101 in additional detail. When training, models 101 receive the input audio stream 210 and a corresponding transcription 211 of the input audio stream. Transcription 211 may be meticulously generated by a human based on the input audio stream 210. Transcription 211 may be converted into a stream of phonemes 213 by system dictionary 212. System dictionary 212 includes correspondences between the written orthographic representation of a word and the phonemes that correspond to the word. A phoneme is generally defined as the smallest acoustic event that distinguishes one word from another.

[0007] Models 101 may model each individual phoneme in the audio stream in the context of its surrounding phonemes. Thus, model 101 may model a target sound using the phonemes that immediately proceed and succeed the target sound. A more complex model may use two proceeding and two succeeding phonemes. In English, there are approximately 50 phonemes. Accordingly, for a model using one phoneme on each side of the target phoneme, there are approximately 125,000 (50^3) possible phoneme groups to model. For a model using two phonemes on each side, there are 322,500,000 (50^5) possible

phoneme groups. Because of the large number of phoneme groups, individually modeling each of the phoneme groups could require impractically complex models. Additionally, a training set that fully covered each possible phoneme group would be inordinately large.

[0008] Accordingly, conventional speech recognition models may be based on phoneme clusters in which each phoneme cluster corresponds to a number of individual phoneme groups. One phoneme cluster, for example, may correspond to all phoneme groups in which the middle sound is a hard “C” sound. A model may use only 2,000-5,000 phoneme clusters instead of the 125,000 or more possible phoneme groups. The reduced number of phoneme clusters relative to the original phoneme groups allows for less complex speech recognition models and smaller training data sets.

[0009] One problem associated with phoneme clustering is determining the mapping of which phoneme groups to assign to which phoneme clusters. One conventional technique for mapping groups to clusters is based on a statistical analysis of the training data. A problem with this technique is that it does not generalize well to phoneme groups that are not present in the training data.

[0010] A second conventional technique for mapping groups to clusters is based on a clustering decision tree. A decision tree uses a predetermined series of questions to categorize a phoneme group into one of the phoneme clusters. For example, a decision tree may first determine whether the left-most phoneme in a group corresponds to a vowel sound. If so, a second question, such as whether the right-most phoneme corresponds to a fricative sound, is asked. If

not, a different second question, such as whether the left-most phoneme corresponds to a silence, is asked. Eventually, the answer to a question will lead to a node of the decision tree that assigns the phoneme grouping to a cluster. In this manner, each phoneme group is assigned to a cluster based on a sequence of questions where the particular sequence is determined by nodes in the clustering tree.

[0011] In contrast to mapping clusters based on a statistical analysis of the training data, classifying phoneme groups based on a decision tree tends to lead to more accurate clustering when classifying data not in the training data.

Phoneme groups that were not in the training data but that are similar to groups in the training data will tend to follow the same path down the clustering tree as the similar phoneme groups. Accordingly, the clustering tree generalizes well when dealing with new phoneme groups.

[0012] A problem associated with clustering trees is that the questions in the trees and the topology of the trees are manually designed by a speech expert. A tree can be complex and is often designed for a specific speech recognition model. Speech recognition systems often include multiple models that, therefore, require the separate creation and maintenance of multiple clustering trees. This can be costly and the independent nature of the clustering trees and the models can lead to non-optimal overlapping results.

[0013] Thus, there is a need in the art for improved clustering trees.

SUMMARY OF THE INVENTION

[0014] Systems and methods consistent with the present invention implement speech recognition using a unified clustering tree from which multiple speech recognition models receive input phoneme clusters. The unified clustering tree is segmented into levels from which the models are derived.

[0015] One aspect of the invention is directed to a data structure encoded on a computer readable medium. The data structure includes data that defines decision nodes associated with questions that relate to a group of sounds and data that defines terminal nodes that define a sound cluster to which the group of sounds belong. The decision nodes and the terminal nodes are defined hierarchically relative to one another and the decision nodes and the terminal nodes are divided into levels, where each level is associated with a different speech recognition model.

[0016] A second aspect consistent with the invention is directed to a method of building a clustering tree for classifying a group of sounds into one of a number of possible sound clusters. The method includes building a first level of the clustering tree with a first hierarchical arrangement of decision nodes in which each decision node of the first hierarchical arrangement is associated with one of a first group of questions relating to the group of sounds. The method further includes building a second level of the clustering tree with a second hierarchical arrangement of decision nodes in which each decision node of the second hierarchical arrangement is associated with one of a second group of questions relating to the group of sounds. The second group of questions

discriminating at a finer level of granularity within the group of sounds than the first group of questions.

[0017] A third aspect consistent with the invention is directed to a speech recognition system that comprises a clustering tree and speech recognition models. The clustering tree classifies a series of sounds into predefined clusters based on one of the sounds and on a predetermined number of neighboring sounds that surround the one of the sounds. The speech recognition models are trained to recognize speech based on the predefined clusters, where each of the speech recognition models receives the predefined clusters from a different portion of the clustering tree.

BRIEF DESCRIPTION OF THE DRAWINGS

[0018] The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate the invention and, together with the description, explain the invention. In the drawings,

[0019] Fig. 1 is a diagram illustrating a conventional speech recognition system;

[0020] Fig. 2 is a diagram illustrating training of models of the speech recognition system of Fig. 1;

[0021] Fig. 3 is a diagram illustrating an exemplary system in which concepts consistent with the invention may be implemented;

[0022] Fig. 4 is a diagram illustrating training of speech recognition models;

[0023] Fig. 5 is a diagram illustrating a unified clustering tree consistent with an aspect of the invention for generating phoneme clusters; and

[0024] Fig. 6 is a flow chart illustrating building of the unified clustering tree shown in Fig. 5.

DETAILED DESCRIPTION

[0025] The following detailed description of the invention refers to the accompanying drawings. The same reference numbers may be used in different drawings to identify the same or similar elements. Also, the following detailed description does not limit the invention. Instead, the scope of the invention is defined by the appended claims and equivalents of the claim limitations.

[0026] Systems and methods consistent with the present invention include a unified clustering tree that generates phoneme clusters for multiple speech recognition models. The unified clustering tree is designed such that different levels of the tree generate the phoneme clusters for the different models. The different models may include different model classes, such as triphone models and quinphone models.

SYSTEM OVERVIEW

[0027] Speech recognition, as described herein, may be performed on one or more processing devices or networks of processing devices. Fig. 3 is a diagram illustrating an exemplary system 300 in which concepts consistent with the invention may be implemented. System 300 includes a computing device 301

that has a computer-readable medium 309, such as random access memory, coupled to a processor 308. Computing device 301 may also include a number of additional external or internal devices. An external input device 320 and an external output device 321 are shown in Fig. 3. Input devices 320 may include, without limitation, a mouse, a CD-ROM, and/or a keyboard. Output devices 321 may include, without limitation, a display and/or an audio output device, such as a speaker.

[0028] In general, computing device 301 may be any type of computing platform, and may be connected to a network 302. Computing device 301 is exemplary only. Concepts consistent with the present invention can be implemented on any computing device, whether or not connected to a network.

[0029] Processor 308 executes program instructions stored in memory 309. Processor 308 can be any of a number of well-known computer processors, such as processors from Intel Corporation, of Santa Clara, California.

[0030] Memory 309 may contain application programs and data. In particular, memory 309 may include a speech recognition component 315. Speech recognition component 315 may contain HMM-based speech recognition models 316. Models 316 may be trained using phoneme clusters assigned by a clustering tree.

SYSTEM OPERATION

[0031] Fig. 4 is a diagram illustrating training of speech recognition models 316 according to an implementation consistent with the invention. Models 316

may be implemented as a hierarchical arrangement of three separate models, labeled as coarse model 401, middle model 402, and fine model 403.

[0032] During training, coarse model 401 receives the training audio stream and an indication of the phoneme clusters that correspond to the audio stream. The phoneme cluster input to coarse model 401 includes phoneme clusters, labeled as level 1 phoneme clusters, generated by a first level of a unified cluster tree (described in more detail below). Similarly, the phoneme cluster input to middle model 402 includes phoneme clusters, labeled as level 2 phoneme clusters, generated by a second level of the unified cluster tree, and the phoneme cluster input to fine model 403 includes phoneme clusters, labeled as level 3 phoneme clusters, generated by a third level of the unified cluster tree. Middle model 402 may also receive the level 1 phoneme clusters and fine model 403 may also receive the level 1 and level 2 phoneme clusters.

[0033] Coarse model 401 is designed such that, during run-time operation, coarse model 401 performs a rough recognition of the speech in the incoming audio stream. A goal of coarse model 401 is not to omit correct answers. In general, finer speech models have a greater risk of incorrectly recognizing speech. Coarse model 401 performs a “first pass” of the input speech by minimizing risk at the expense of a finer grain recognition. Coarse model 401 may include a triphone non-crossword model. A triphone model refers to a language model that uses three phonemes — two context phonemes that surround the phoneme that is being modeled — when making phoneme

recognition decisions. Non-crossword means that the model does not attempt to deal with phonemes that cross word boundaries.

[0034] Middle model 402 refines the speech recognized by coarse model 401. Middle model 402 refines the output of coarse model 401 by more precisely recognizing phonemes in the speech. The phonemes recognized by middle model 402 have a higher risk of incorrect classification than those recognized by coarse model 401. Middle model 402 may include a quinphone non-crossword model. A quinphone model uses five phonemes — four context phonemes that surround the phoneme that is being modeled — when making phoneme recognition decisions.

[0035] Fine model 403 refines the classifications from middle model 402. Fine model 403 is designed to recognize phonemes in the incoming audio stream to a high level of detail. Fine model 403 may include a quinphone crossword model.

[0036] Models 401-403 may be trained on phoneme clusters instead of directly on the triphone and quinphone phoneme groups. Fig. 5 is a diagram illustrating a unified clustering tree 500 consistent with an aspect of the invention for generating the phoneme clusters. Clustering tree 500 can be conceptually broken into three levels: a first level 501, a second level 502, and a third level 503. The clusters generated by first level 501 (level 1 phoneme cluster) are used by coarse model 401. Similarly, the clusters generated by second level 502 (level 2 phoneme cluster) and third level 503 (level 3 phoneme cluster) are used by middle model 402 and fine model 403, respectively. Middle model 402 may

additionally use the clusters generated for coarse model 401 and fine model 403 may additionally use the clusters generated for coarse model 401 and middle model 402.

[0037] Unified clustering tree 500 includes a number of decision nodes, such as nodes 510a – 510c (shown as darkened circles), and terminal nodes, such as nodes 511a – 511b (shown as open circles). Decision nodes 510 correspond to questions or decision criteria that is applied to the current phoneme group. The response at a decision node 510 is used by speech recognition component 315 in determining how to move down to the next node in clustering tree 500. This process is repeated until a terminal node is encountered.

[0038] Terminal nodes 511 correspond to a single phoneme cluster. Thus, terminal node 511a may correspond to a first cluster and terminal node 511b may correspond to a second cluster. Each phoneme group applied to clustering tree 500 is classified into the phoneme cluster based on the terminal node 511 in which the phoneme group is placed.

[0039] For ease of explanation, unified clustering tree 500 is a smaller clustering tree than that which may be typically implemented. A more useful unified clustering tree may have, for example, upwards of 2,000 nodes.

[0040] As an example of traversing unified clustering tree 500, consider the phoneme group “D-*-C-A-T,” where “*” indicates silence. The center phoneme “C” is the phoneme being modeled. Node 510a may determine whether the phoneme to the immediate left of the center phoneme represents a vowel sound. In this example, the answer is no, and the clustering tree proceeds to node 510b.

Node 510b may determine whether the phoneme to the left of the center phoneme represents silence. In this example, the answer is yes, and the decision tree proceeds to node 510c. In this manner, unified clustering tree is traversed until a terminal node is encountered. The terminal node indicates to which cluster the phoneme group is assigned.

[0041] The specific questions to use for each of nodes 510 may be determined by a speech expert based on the particular speech recognition problem. In general, unified clustering tree 500 is designed such that the questions become more detailed as the clustering tree progresses from level 501 to level 503. Questions in level 501, for example, may be based only on the center phoneme and its two surrounding phonemes while questions in levels 502 and 503 may be based on the center phoneme and its four surrounding phonemes. The triphone coarse model 401, which is based on the clusters identified by level 501, follows from the questions corresponding to the nodes in level 501. Similarly, quinphone models 402 and 403 follow from the questions corresponding to the nodes in levels 502 and 503. In this manner, unified clustering tree 500 drives the creation of models 401-403.

[0042] Fig. 6 is a flow chart illustrating building of a unified clustering tree, such as unified clustering tree 500, consistent with an aspect of the invention. A designer, such as a speech expert, builds level 501 of the clustering tree using "coarse" phoneme classification questions (Act 601). The questions in level 501 may be based on the center phoneme and its two surrounding context phonemes. Questions may additionally be based on externally measurable

features, such as whether music is present. The particular questions to ask and the number of nodes to use generally depend on the speech recognition problem being addressed and can be influenced by factors, such as language that is to be recognized. Building a single level clustering tree is generally known in the art.

[0043] One method of determining the effectiveness of a level is based on the entropy of the level. Accordingly, the designer may use the entropy of a level as feedback relating to the effectiveness of the level. In particular, the designer may calculate the entropy for the first level and accept the first level when the entropy is below a preset threshold (Acts 602 and 603).

[0044] At this point, the first level may be frozen and not further modified (Act 604).

[0045] The designer then builds the second level starting from the end nodes of the first level (Act 605). The entropy for the second level is used as feedback to determine when the second level is complete (Acts 606 and 607). When it is complete, the second level may be frozen (Act 608).

[0046] At this point, the designer may then build the third level starting from the end nodes of the second level (Act 609). The entropy for the third level is used as feedback to determine when the third level is complete (Acts 610 and 611). When it is complete, the third level may be frozen (Act 612).

[0047] The process shown in Fig. 6 may be repeated for any additional levels that are desired. Although unified clustering tree 500 was shown having three levels, the unified clustering tree could be designed using the above-described concepts with two levels or more than two levels.

[0048] The type of questions corresponding to nodes 510 in each unified clustering tree will now be described in further detail.

CONCLUSION

[0049] As described above, a unified clustering tree includes a number of levels that generate phoneme clusters. A different type of speech recognition model can be derived from each level. Models at higher levels in the unified clustering tree are coarse models relative to more fine-grain models at lower levels of the clustering tree.

[0050] The foregoing description of preferred embodiments of the invention provides illustration and description, but is not intended to be exhaustive or to limit the invention to the precise form disclosed. Modifications and variations are possible in light of the above teachings or may be acquired from practice of the invention. For example, while a series of acts has been presented with respect to Fig. 6, the order of the acts may be different in other implementations consistent with the present invention. Additionally, non-dependent acts may be implemented in parallel.

[0051] Certain portions of the invention have been described as software that performs one or more functions. The software may more generally be implemented as any type of logic. This logic may include hardware, such as an application specific integrated circuit a field programmable gate array, software, or a combination of hardware and software.

[0052] No element, act, or instruction used in the description of the present application should be construed as critical or essential to the invention unless explicitly described as such. Also, as used herein, the article “a” is intended to include one or more items. Where only one item is intended, the term “one” or similar language is used.

[0053] The scope of the invention is defined by the claims and their equivalents.